

The Complexity of Grid Coloring

Daniel Apon *

Univ. of MD at College Park

William Gasarch †

Univ. of MD at College Park

Kevin Lawler ‡

Permanent

March 4, 2013

Abstract

A c -coloring of $G_{n,m} = [n] \times [m]$ is a mapping of $G_{n,m}$ into $[c]$ such that no four corners forming a rectangle have the same color. In 2009 a challenge was proposed via the internet to find a 4-coloring of $G_{17,17}$. This attracted considerable attention from the popular mathematics community. A coloring was produced; however, finding it proved to be difficult. The question arises: is the problem of grid coloring is difficult in general? We present three results that support this conjecture: (1) Given a partial c -coloring of an $G_{n,m}$ grid, can it be extended to a full c -coloring? We show this problem is NP-complete. (2) The statement $G_{n,m}$ is c -colorable can be expressed as a Boolean formula with nmc variables. We show that if the $G_{n,m}$ is not c -colorable then any tree resolution proof of the corresponding formula is of size $2^{\Omega(c)}$. (We then generalize this result for other monochromatic shapes.) (3) We show that any tree-like cutting planes proof that $c+1$ by $c\binom{c+1}{2}+1$ is not c -colorable must be of size $2^{\Omega(c^3/\log^2 c)}$. Note that items (2) and (3) yield statements from Ramsey Theory which are of size polynomial in their parameters and require exponential size in various proof systems.

1 Introduction

Notation 1.1 If $x \in \mathbb{N}$ then $[x]$ denotes the set $\{1, \dots, x\}$. $G_{n,m}$ is the set $[n] \times [m]$. If X is a set and $k \in \mathbb{N}$ then $\binom{X}{k}$ is the set of all size- k subsets of X .

On November 30, 2009 the following challenge was posted on Complexity Blog [3].

BEGIN EXCERPT

The 17×17 challenge: worth \$289.00. I am not kidding.

*University of Maryland, College Park, MD 20742. dapon@cs.umd.edu

†University of Maryland, College Park, MD 20742. gasarch@cs.umd.edu

‡Permanent, Berkeley, CA 94710. kevin@permanentco.com

Def 1.2 A *rectangle* of $G_{n,m}$ is a subset of the form $\{(a, b), (a+c_1, b), (a+c_1, b+c_2), (a, b+c_2)\}$ for some $a, b, c_1, c_2 \in \mathbb{N}$. A grid $G_{n,m}$ is *c-colorable* if there is a function $\chi : G_{n,m} \rightarrow [c]$ such that there are no rectangles with all four corners the same color.

The 17×17 challenge: The first person to email me a 4-coloring of $G_{17,17}$ in LaTeX will win \$289.00. (289.00 is chosen since it is 17^2 .)

END EXCERPT

There are two motivations for this kind of problem. (1) The problem of coloring grids to avoid rectangles is a relaxation of the classic theorem (a corollary of the Gallai-Witt theorem) which states that for a large enough grid any coloring yields a monochromatic square, and (2) grid-coloring problems avoiding rectangles are equivalent to finding certain bipartite Ramsey Numbers. For more details on these motivations, and why the four-coloring of $G_{17,17}$ was of particular interest, see the post [3]. or the paper by Fenner, et al [2].

Brian Hayes, the Mathematics columnist for Scientific American, publicized the challenge [6]. Initially there was a lot of activity on the problem. Some used SAT solvers, some used linear programming, and one person offered an exchange: *buy me a \$5000 computer and I'll solve it*. Finally in 2012 Bernd Steinbach and Christian Posthoff [15, 5] solved the problem. They used a rather clever algorithm with a SAT solver. They believed that the solution was close to the limits of their techniques.

Though this particular instance of the problem was solved, the problem of grid coloring in general seems to be difficult. In this paper we formalize and prove three different results that indicate grid coloring is hard.

1.1 Grid Coloring Extension is NP-Complete

Between the problem being posed and resolved the following challenge was posted [4] though with no cash prize. We paraphrase the post.

BEGIN PARAPHRASE

Def 1.3 Let $c, N, M \in \mathbb{N}$.

1. A mapping χ of $G_{N,M}$ to $[c]$ is a *c-coloring* if there are no monochromatic rectangles.
2. A partial mapping χ of $G_{N,M}$ to $[c]$ is *extendable to a c-coloring* if there is an extension of χ to a total mapping which is a *c-coloring* of $G_{N,M}$. We will use the term *extendable* if the c is understood.

Def 1.4 Let

$$GCE = \{(N, M, c, \chi) \mid \chi \text{ is extendable}\}.$$

GCE stands for *Grid Coloring Extension*.

CHALLENGE: Prove that GCE is NP-complete.
END PARAPHRASE

In Section 2 we show that GCE is indeed NP-complete. This result may explain why the original 17×17 challenge was so difficult. Then again—it may not. In Section 4 we show that GCE is fixed-parameter tractable. Hence, for a fixed c , the problem might not be hard. In Section 8 we state open problems.

There is another reason the results obtained may not be the reason why the 17×17 challenge was hard. The 17×17 challenge can be rephrased as proving that $(17, 17, 4, \chi) \in GCE$ where χ is the empty partial coloring. This is a rather special case of GCE since *none* of the spots are pre-colored. It is possible that GCE in the special case where χ is the empty coloring is easy. While we doubt this is true, we note that we have not eliminated the possibility.

One could ask about the problem

$$GC = \{(n, m, c) \mid G_{n,m} \text{ is } c\text{-colorable}\}.$$

However, this does not quite work. If n, m are in unary, then GC is a sparse set. By Mahaney's Theorem [12, 7] if a sparse set is NP-complete then $P = NP$. If n, m are in binary, then we cannot show that GC is in NP since the obvious witness is exponential in the input. This formulation does not get at the heart of the problem, since we believe it is hard because the number of possible colorings is large, not because n, m are large. It is an open problem to find a framework within which a problem like GC can be shown to be hard.

1.2 Grid Coloring is Hard for Tree Res

The statement $G_{n,m}$ is c -colorable can be written as a Boolean formula (see Section 5). If $G_{n,m}$ is not c -colorable then this statement is not satisfiable. A *Resolution Proof* is a formal type of proof that a formula is not satisfiable. One restriction of this is *Tree Resolution*.

In Section 5 we define all of these terms. We then show that any tree resolution of the Boolean Formula representing $G_{n,m}$ is c -colorable requires size $2^{\Omega(c)}$.

1.3 A Particular Grid Coloring Problem is Hard for Tree-Like Cutting Planes Proofs

The statement $G_{n,m}$ is c -colorable is equivalent to the statement $A\vec{x} \leq \vec{b}$ has no 0-1 solution for some matrix A and vector \vec{b} . (Written as $A\vec{x} \leq \vec{b} \notin SAT$.) It is known [2] that $G_{n,m}$ is not c -colorable when $n = c + 1$ and $m = c \binom{c}{2} + 1$. A *Cutting Planes Proof* is a formal type of proof that $A\vec{x} \leq \vec{b} \notin SAT$. One restriction of this is *Tree-like Cutting Plane Proofs*.

In Section 7 we define all of these terms. We then show that any tree-like CP proof of $A\vec{x} \leq \vec{b} \notin SAT$, where this is equivalent to $G_{c+1, c \binom{c}{2} + 1}$ not being c -colorable, requires size $2^{\Omega(c^3 / \log^2 c)}$.

This lower bound on tree-like CP proofs yields a lower bound on tree-res proofs of the statement that $G_{c+1, c\binom{c}{2}+1}$ is not c -colorable of $2^{\Omega(c^3/\log^2 c)}$. This is not too far away from the upper bound of $O(c^4)$.

2 GCE is NP-complete

Theorem 2.1 *GCE is NP-complete.*

Proof:

Clearly $GCE \in \text{NP}$.

Let $\phi(x_1, \dots, x_n) = C_1 \wedge \dots \wedge C_m$ be a 3-CNF formula. We determine N, M, c and a partial c -coloring χ of $G_{N,M}$ such that

$$(N, M, c, \chi) \in GCE \text{ iff } \phi \in 3\text{-SAT}.$$

The grid will be thought of as a main grid with irrelevant entries at the left side and below, which are only there to enforce that some of the colors in the main grid occur only once. The colors will be T, F , and some of the $(i, j) \in G_{N,M}$. We use (i, j) to denote a color for a particular position.

The construction is in four parts. We summarize the four parts here before going into details.

1. We will often need to define $\chi(i, j)$ to be (i, j) and then never have the color (i, j) appear in any other cell of the main grid. We show how to color the cells that are not in the main grid to achieve this. While we show this first, it is actually the last step of the construction.
2. The main grid will have $2nm + 1$ rows. In the first column we have $2nm$ blank spaces and the space $(1, 2nm + 1)$ colored with $(1, 2nm + 1)$. The $2nm$ blank spaces will be forced to be colored T or F . We think of the column as being in n blocks of $2m$ spaces each. In the i th block the coloring will be forced to be

$$\begin{array}{c} T \\ F \\ \vdots \\ T \\ F \end{array}$$

if x_i is to be set to T , or

$$\begin{array}{c} F \\ T \\ \vdots \\ F \\ T \end{array}$$

if x_i is to be set to F .

3. For each clause C there will be two columns. The coloring χ will be defined on most of the cells in these columns. However, the coloring will extend to these two columns iff one of the literals in C is colored T in the first column.
4. We set the number of colors properly so that the T and F will be forced to be used in all blank spaces.

1) **Forcing a color to appear only once in the main grid.**

Say we want the cell $(2, 4)$ in the main grid to be colored $(2, 4)$ and we do not want this color appearing anywhere else in the main grid. We can do the following: add a column of $(2, 4)$'s to the left end (with one exception) and a row of $(2, 4)$'s below. Here is what we get:

(2, 4)								
(2, 4)								
T		(2, 4)						
(2, 4)								
(2, 4)								
(2, 4)								
(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)

(The double lines are not part of the construction. They are there to separate the main grid from the rest.)

It is easy to see that in any coloring of the above grid the only cells that can have the color $(2, 4)$ are those shown to already have that color. It is also easy to see that the color T we have will not help to create any monochromatic rectangles since there are no other T 's in its column. The T we are using *is* the same T that will later mean TRUE. We could have used F . If we used a new special color we would need to be concerned whether there is a monochromatic grid of that color. Hence we use T .

What if some other cell needs to have a unique color? Lets say we also want to color cell $(5, 3)$ in the main grid with $(5, 3)$ and do not want to color anything else in the main grid $(5, 3)$. Then we do the following:

(5, 3)	(2, 4)								
(5, 3)	(2, 4)								
(5, 3)	T		(2, 4)						
T	(2, 4)					(5, 3)			
(5, 3)	(2, 4)								
(5, 3)	(2, 4)								
(5, 3)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 4)
(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)	(5, 3)

It is easy to see that in any coloring of the above grid the only cells that can have the color $(2, 4)$ or $(5, 3)$ are those shown to already have those colors.

For the rest of the construction we will only show the main grid. If we denote a color as D (short for *Distinct*) in the cell (i, j) then this means that (1) cell (i, j) is color (i, j) and (2) we have used the above gadget to make sure that (i, j) does not occur as a color in any other cell of the main grid. Note that we when we have D in the $(2, 4)$ cell and in the $(5, 3)$ cell they denote different colors.

2) **Forcing (x, \bar{x}) to be colored (T, F) or (F, T) .**

There will be one column with cells labeled by literals. The cells are blank, uncolored. We will call this row *the literal column*. We will put to the left of the literal column, separated by a triple line, the literals whose values we intend to set. These literals are not part of the construction; they are a visual aid. The color of the literal-labeled cells will be T or F . We need to make sure that all of the x_i have the same color and that the color is different than that of \bar{x}_i .

Here is an example which shows how we can force (x_1, \bar{x}_1) to be colored (T, F) or (F, T) .

\bar{x}_1		T	F
x_1		T	F

We will actually need m copies of x_1 and m copies of \bar{x}_1 . We will also put a row of D 's on top which we will use later. We illustrate how to do this in the case of $m = 3$.

	D	D	D	D	D	D	D	D	D	D	D
\bar{x}_1		D	D	D	D	D	D	D	D	T	F
x_1		D	D	D	D	D	T	F	T	F	F
\bar{x}_1		D	D	D	D	T	F	T	F	D	D
x_1		D	D	T	F	T	F	D	D	D	D
\bar{x}_1		T	F	T	F	D	D	D	D	D	D
x_1		T	F	D	D	D	D	D	D	D	D

We leave it as an exercise to prove that

- If the bottom x_1 cell is colored T then (1) all of the x_1 cells are colored T , and (2) all of the \bar{x}_1 cells are colored F .
- If the bottom x_1 cell is colored F then (1) all of the x_1 cells are colored F , and (2) all of the \bar{x}_1 cells are colored T .

Note that (1) if we want one literal-pair (that is x_1, \bar{x}_1) then we use two columns, (2) if we want two literal-pairs then we use six columns, and (3) if we want three literal-pairs then we use ten columns. We leave it as an exercise to generalize the construction to m literal-pairs using $2 + 4(m - 1)$ columns.

We will need m copies of x_2 and m copies of \bar{x}_2 . We illustrate how to do this in the case of $m = 2$. We use double lines in the picture to clarify that the x_1 and the x_2 variables are not chained together in any way.

	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D
\bar{x}_2		D	D	D	D	D	D	D	D	D	D	T	F		
x_2		D	D	D	D	D	D	D	D	T	F	T	F		
\bar{x}_2		D	D	D	D	D	D	D	D	T	F	T	F	D	D
x_2		D	D	D	D	D	D	D	D	T	F	D	D	D	D
\bar{x}_1		D	D	D	D	T	F	D	D	D	D	D	D	D	D
x_1		D	D	T	F	T	F	D	D	D	D	D	D	D	D
\bar{x}_1		T	F	T	F	D	D	D	D	D	D	D	D	D	D
x_1		T	F	D	D	D	D	D	D	D	D	D	D	D	D

We leave it as an exercise to prove that, for all $i \in \{1, 2\}$:

- If the bottom x_i cell is colored T then (1) all of the x_i cells are colored T , and (2) all of the \bar{x}_1 cells are colored F .
- If the bottom x_i cell is colored F then (1) all of the x_i cells are colored F , and (2) all of the \bar{x}_1 cells are colored T .

An easy exercise for the reader is to generalize the above to a construction with n variables with m literal-pairs for each variable. This will take $n(2 + 4(m - 1))$ columns.

For the rest of the construction we will only show the literal column and the clause columns (which we define in the next part). It will be assumed that the D 's and T 's and F 's are in place to ensure that all of the x_i cells are one of $\{T, F\}$ and the \bar{x}_i cells are the other color.

3) How we force the coloring to satisfy ONE clause

Say one of the clauses is $C_1 = L_1 \vee L_2 \vee L_3$ where L_1, L_2 , and L_3 are literals. Pick an L_1 row, an L_2 row, and an L_3 row. We will also use the top row, as we will see. For other clauses you will pick other rows. Since there are m copies of each variable and its negation this is easy to do.

The two T 's in the top row in the next picture are actually in the very top row of the grid.

We put a C_1 over the columns that will enforce that C_1 is satisfied. We put L_1, L_2 , and L_3 on the side to indicate the positions of the variables. These C_1 and the L_i outside the triple bars are not part of the grid. They are a visual aid.

		C_1	C_1
	D	T	T
L_3		D	F
L_2			
L_1		F	D

Claim 1: If χ' is a 2-coloring of the blank spots in this grid (with colors T and F) then it CANNOT have the L_1, L_2, L_3 spots all colored F .

Proof of Claim 1:

Assume, by way of contradiction, that that L_1, L_2, L_3 are all colored F . Then this is what it looks like:

		C_1	C_1
	D	T	T
L_3	F	D	F
L_2	F		
L_1	F	F	D

The two blank spaces are both FORCED to be T since otherwise you get a monochromatic rectangle of color F . Hence we have

		C_1	C_1
	D	T	T
L_3	F	D	F
L_2	F	T	T
L_1	F	F	D

This coloring has a monochromatic rectangle which is colored T . This contradicts χ' being a 2-coloring of the blank spots.

End of Proof of Claim 1

We leave the proof of Claim 2 below to the reader.

Claim 2: If χ' colors L_1, L_2, L_3 anything except F, F, F then χ' can be extended to a coloring of the grid shown.

Upshot: A 2-coloring of the grid is equivalent to a satisfying assignment of the clause.

Note that each clause will require 2 columns to deal with. So there will be $2m$ columns for this.

4) Putting it all together

Recall that $\phi(x_1, \dots, x_n) = C_1 \wedge \dots \wedge C_m$ is a 3-CNF formula.

We first define the main grid and later define the entire grid and N, M, c .

The main grid will have $2nm + 1$ rows and $n(4m - 2) + 2m + 1$ columns. The first $n(4m - 2) + 1$ columns are partially colored using the construction in Part 2. This will establish the literal column. We will later set the number of colors so that the literal column must use the colors T and F .

For each of the m clauses we pick a set of its literals from the literals column. These sets-of-literals are all disjoint. We can do this since we have m copies of each literal-pair. We then do the construction in Part 3. Note that this uses two columns. Assuming that all of the D 's are colored distinctly and that the only colors left are T and F , this will ensure that the main grid is c -colorable iff the formula is satisfiable.

The main grid is now complete. For every (i, j) that is colored (i, j) we perform the method in Part 1 to make sure that (i, j) is the only cell with color (i, j) . Let the number of such (i, j) be C . The number of colors c is $C + 2$. ■

3 An Example

We can make the construction slightly more efficient (and thus can actually work out an example). We took m pairs $\{x_i, \bar{x}_i\}$. We don't really need all m . If x_i appears in a clauses and \bar{x}_i appears in b clauses then we only need $\max\{a, b\}$ literal-pairs. If $a \neq b$ then we only need $\max\{a, b\} - 1$ literal-pairs and one additional literal. (This will be the case in the example below.)

With this in mind we will do an example- though we will only show the main grid.

$$(x_1 \vee x_2 \vee \bar{x}_3) \wedge (\bar{x}_2 \vee x_3 \vee x_4) \wedge (\bar{x}_1 \vee \bar{x}_3 \vee \bar{x}_4)$$

We only need

- one (x_1, \bar{x}_1) literal-pair,
- one (x_2, \bar{x}_2) literal-pair,
- one (x_3, \bar{x}_3) literal-pair,
- one additional \bar{x}_3 ,
- one (x_4, \bar{x}_4) literal-pair.

													C_1	C_1	C_2	C_2	C_3	C_3
	D	D	D	D	D	D	D	D	D	D	D	D	T	T	T	T	T	T
\bar{x}_4		D	D	D	D	D	D	D	D	T	F	D	D	D	D	D	D	F
x_4		D	D	D	D	D	D	D	D	T	F	D	D	D	F	D	D	D
\bar{x}_3		D	D	D	D	D	D	T	F	D	D	D	D	D	D	D	D	D
x_3		D	D	D	D	T	F	T	F	D	D	D	D			D	D	
\bar{x}_3		D	D	D	D	T	F	D	D	D	D	D	F	D	D			
\bar{x}_2		D	D	T	F	D	D	D	D	D	D	D	D	F	D	D	D	D
x_2		D	D	T	F	D	D	D	D	D				D	D	D	D	D
\bar{x}_1		T	F	D	D	D	D	D	D	D	D	D	D	D	D	D	F	D
x_1		T	F	D	D	D	D	D	D	D	D	F	D	D	D	D	D	D

4 Fixed Parameter Tractability

The 17×17 problem only involved 4-colorability. Does the result that GCE is NP-complete really shed light on the hardness of the 17×17 problem? What happens if the number of colors is fixed?

Def 4.1 Let $c \in \mathbb{N}$. Let

$$GCE_c = \{(N, M, \chi) \mid \chi \text{ can be extended to a } c\text{-coloring of } G_{N,M}\}.$$

Clearly $GCE_c \in DTIME(c^{O(NM)})$. Can we do better? Yes. We will show that GCE is in time $O(N^2M^2 + 2^{O(c^4)})$.

Lemma 4.2 *Let n, m, c be such that $c \leq 2^{nm}$. Let χ be a partial c -coloring of $G_{n,m}$. Let U be the uncolored grid points. Let $|U| = u$. There is an algorithm that will determine if χ can be extended to a full c -coloring that runs in time $O(cnm2^{2u}) = 2^{O(nm)}$.*

Proof: For $S \subseteq U$ and $1 \leq i \leq c$ let

$$f(S, i) = \begin{cases} YES & \text{if } \chi \text{ can be extended to color } S \text{ using only colors } \{1, \dots, i\}; \\ NO & \text{if not.} \end{cases} \quad (1)$$

We assume throughout that the coloring χ has already been applied.

We are interested in $f(U, c)$; however, we use a dynamic program to compute $f(S, i)$ for all $S \subseteq U$ and $1 \leq i \leq c$. Note that $f(\emptyset, i) = YES$.

We describe how to compute $f(S, i)$. Assume that for all S' such that $|S'| < |S|$, for all $1 \leq i \leq c$, $f(S', i)$ is known.

1. For all nonempty 1-colorable $T \subseteq S$ do the following (Note that there are at most 2^u sets T .)
 - (a) If $f(S - T, i) = NO$ then $f(S, i) = NO$.
 - (b) If $f(S - T, i - 1) = YES$ then determine if coloring T with i will create a monochromatic rectangle. If not then $f(S, i) = YES$. Note that this takes $O(nm)$.
2. We now know that for all 1-colorable $T \subseteq S$ (1) $f(S - T, i) = YES$, and (2) either $f(S - T, i - 1) = NO$ or $f(S - T, i - 1) = YES$ and coloring T with i creates a monochromatic rectangle. We will show that in this case $f(S, i) = NO$.

Assume that, for all 1-colorable sets $T \subseteq S$: (1) $f(S - T, i) = YES$, and (2) either $f(S - T, i - 1) = NO$ or $f(S - T, i - 1) = YES$ and coloring T with i creates a rectangle with χ . Also assume, by way of contradiction, that $f(S, i) = YES$. Let COL be an extension of χ to S . Let T be the set colored i . Clearly $f(S - T, i - 1) = YES$. Hence the second clause of condition (2) must hold. Hence coloring T with i creates a monochromatic rectangle. This contradicts COL being a c -coloring.

The dynamic program fills in a table that is indexed by the 2^u subsets of S and the c colors. Each slot in the table takes $O(nm2^u)$ to compute. Hence to fill the entire table takes $O(cnm2^{2u})$ steps. ■

Lemma 4.3 *Assume $c + 1 \leq N$ and $c \binom{c+1}{2} < M$. Then $G_{N,M}$ is not c -colorable. Hence, for any χ , $(N, M, \chi) \notin GCE_c$.*

Proof: Assume, by way of contradiction, that there is a c -coloring of $G_{N,M}$. Since every column has at least $c + 1$ elements the following mapping is well defined: Map every column to the least $(\{i, j\}, a)$ such that the $\{i, j\} \in \binom{[c+1]}{2}$ and both the i th and the j th row of that column are colored a . The range of this function has $c \binom{c+1}{2}$ elements. Hence some element of the range is mapped to at least twice. This yields a monochromatic rectangle. ■

Lemma 4.4 *Assume $N \leq c$ and $M \in \mathbb{N}$. If χ is a partial c -coloring of $G_{N,M}$ then $(N, M, \chi) \in GCE_c$.*

Proof: The partial c -coloring χ can be extended to a full c -coloring as follows: for each column use a different color for each blank spot, making sure that all of the new colors in that column are different from each other. ■

Theorem 4.5 $GCE_c \in DTIME(N^2M^2 + 2^{O(c^6)})$ time.

Proof:

1. Input (N, M, χ) .
2. If $N \leq c$ or $M \leq c$ then test if χ is a partial c -coloring of $G_{N,M}$. If so then output YES. If not then output NO. (This works by Lemma 4.4.) This takes time $O(N^2M^2)$. Henceforth we assume $c + 1 \leq N, M$.
3. If $c \binom{c+1}{2} < M$ or $c \binom{c+1}{2} < N$ then output NO and stop. (This works by Lemma 4.3.)
4. The only case left is $c + 1 \leq N, M \leq c \binom{c+1}{2}$. By Lemma 4.2 we can determine if χ can be extended in time $O(2^{NM}) = O(2^{c^6})$.

Step 2 takes $O(N^2M^2)$ and Step 4 takes time $2^{O(c^6)}$, hence the entire algorithm takes time $O(N^2M^2 + 2^{O(c^6)})$. ■

Can we do better? Yes, but it will require a result from [2].

Lemma 4.6 *Let $1 \leq c' \leq c - 1$.*

1. *If $N \geq c + c'$ and $M > \frac{c}{c'} \binom{c+c'}{2}$ then $G_{N,M}$ is not c -colorable.*
2. *If $N \geq 2c$ and $M > 2 \binom{2c}{2}$ then $G_{N,M}$ is not c -colorable. (This follows from a weak version of the $c' = c - 1$ case of Part 1.)*

Theorem 4.7 $GCE_c \in DTIME(N^2M^2 + 2^{O(c^4)})$ time.

Proof:

1. Input (N, M, χ) .
2. If $N \leq c$ or $M \leq c$ then test if χ is a partial c -coloring of $G_{N,M}$. If so then output YES. If not then output NO. (This works by Lemma 4.4.) This takes time $O(N^2M^2)$.
3. For $1 \leq c' \leq c-1$ we have the following pairs of cases.
 - (a) $N = c + c'$ and $M > \frac{c}{c'} \binom{c+c'}{2}$ then output NO and stop. (This works by Lemma 4.6.)
 - (b) $N = c + c'$ and $M \leq \frac{c}{c'} \binom{c+c'}{2}$. By Lemma 4.2 we can determine if χ can be extended to a total c -coloring in time $2^{O(NM)}$. Note that $MN \leq (c + c') \frac{c}{c'} \binom{c+c'}{2}$. On the interval $1 \leq c' \leq c-1$ this function achieves its maximum when $c' = 1$. Hence this case takes $2^{O(c^4)}$.

Henceforth we assume $2c \leq N, M$.

4. If $M > 2 \binom{2c}{2}$ or $N > 2 \binom{2c}{2}$ then output NO and stop. (This works by Lemma 4.6.)
5. The only case left is $2c \leq N, M \leq 2 \binom{2c}{2}$. By Lemma 4.2 we can determine if χ can be extended in time $2^{O(NM)} \leq 2^{O(c^4)}$.

Step 2 and Step 4 together take time $O(N^2M^2 + 2^{O(c^4)})$. ■

Even for small c the additive term $2^{O(c^4)}$ is the real timesink. A cleverer algorithm that reduces this term is desirable. By Theorem 2.1 this term cannot be made polynomial unless $P=NP$.

5 Lower Bound on Tree Res

For n, m, c we define a Boolean formula $GRID(n, m, c)$ such that

$G_{n,m}$ is c -colorable iff $GRID(n, m, c) \in SAT$.

- The variables are x_{ijk} where $1 \leq i \leq n$, $1 \leq j \leq m$, $1 \leq k \leq c$. The intention is that, for all (i, j) , there is a k such that x_{ijk} is true. We interpret k to be the color of (i, j) .
- For all (i, j) we have the clause

$$\bigvee_{k=1}^c x_{ijk}.$$

These clauses ensure that every (i, j) has at least one color.

- For all $1 \leq i < i' \leq n$ and $1 \leq j < j' \leq m$ we have the clause

$$\bigvee_{k=1}^c \neg x_{ijk} \vee \neg x_{i'jk} \vee \neg x_{ij'k} \vee \neg x_{i'j'k}.$$

These clauses ensure there are no monochromatic rectangles.

We do not use clauses to ensure that every (i, j) has at most one color. This is because if the formula above is satisfied then one can extract out of it a c -coloring of $G_{n,m}$ by taking the color of (i, j) to be the *least* k such that x_{ijk} is true.

We show that if $G_{n,m}$ is not c -colorable then any tree resolution proof of $GRID(n, m, c) \notin SAT$ requires size $2^{\Omega(c)}$.

5.1 Background on Tree Resolution and the Prover-Delayer Game

The definitions of Resolution and Tree Resolution are standard. Prover-Delayer games were first defined in [13], however we use the asymmetric version which was first defined in [1]. See also [9].

Def 5.1 Let $\varphi = C_1 \wedge \dots \wedge C_L$ be a CNF formula. A *Resolution Proof that $\varphi \notin SAT$* is a sequence of clauses such that on each line you have either

1. One of the C 's in φ (called an AXIOM).
2. $A \vee B$ where on prior lines you had $A \vee x$ and $B \vee \neg x$.
3. The last line has the empty clause.

It is easy to see that if there is a resolution proof that $\varphi \notin SAT$ then indeed $\varphi \notin SAT$. The converse is also true though slightly harder to prove.

Def 5.2 A *Tree Resolution* proof is one whose underlying graph is a tree.

Def 5.3 The *Prover-Delayer Game* has parameters (1) $a, b \in (1, \infty)$, such that $\frac{1}{a} + \frac{1}{b} = 1$, (2) $p \in \mathbb{R}^+$, and (3) a CNF-formula

$$\varphi = C_1 \wedge \dots \wedge C_L \notin SAT.$$

The game is played as follows until a clause is proven false:

1. The Prover picks a variable x that was not already picked.
2. The Delayer either
 - (a) Sets x to T or F .
 - (b) Defers to the Prover.
 - i. If the Prover sets x to F then the Delayer gets $\lg a$ points.
 - ii. If the Prover sets x to T then the Delayer gets $\lg b$ points.

When some clause has all of its literals set to false the game ends. At that point, if the Delayer has p points then he WINS; otherwise the Prover WINS.

We assume that the Prover and the Delayer play perfectly.

1. *The Prover wins* means *the Prover has a winning strategy*.
2. *The Delayer wins* means *the Delayer has a winning strategy*.

Lemma 5.4 *Let $a, b \in (1, \infty)$ such that $\frac{1}{a} + \frac{1}{b} = 1$, $p \in \mathbb{R}^+$, $\varphi \notin SAT$, φ in CNF-form. If the Delayer wins then EVERY Tree Resolution proof for φ has size $\geq 2^p$.*

Note that the lower bound in Lemma 5.4 is 2^p , not $2^{\Omega(p)}$.

5.2 Lower Bound on Tree Resolution

Theorem 5.5 *Let n, m, c be such that $G_{n,m}$ is not c -colorable and $c \geq 9288$. Any tree resolution proof of $GRID(n, m, c) \notin SAT$ requires size 2^{Dc} where $D = 0.836$.*

Proof:

By Lemma 5.4 it will suffice to show that there exists $a, b \in (1, \infty)$ with $\frac{1}{a} + \frac{1}{b} = 1$, such that the Delayer wins the Prover-Delayer game with parameters a, b, Dc , and $GRID(n, m, c)$. We will determine a, b later. We will also need parameter $r \in (0, 1)$ to be determined.

Here is the Delayers strategy: Assume x_{ijk} was chosen by the Prover.

1. If coloring (i, j) with color k will create a monochromatic rectangle then the Delayer will NOT let this happen—he will set x_{ijk} to F . The Delayer does not get any points but he avoids the game ending. (Formally: if there exists i', j' such that $x_{i'jk} = x_{ijk} = x_{i'j'k} = T$ then the Delayer sets x_{ijk} to F .) Otherwise he goes to the next step of the strategy.
2. If there is a danger that all of the x_{ij*} will be false for some (i, j) then the Delayer will set x_{ijk} to T . The Delayer does not want to panic and set x_{ijk} to T unless he feels he has to. He uses the parameter r . If there are at least rc values k' where the Prover has set $x_{ijk'}$ to F , and there are no $x_{ijk'}$ that have been set to T (by anyone) then Delayer sets x_{ijk} to T . Note that this cannot form a monochromatic rectangle since in step 1 of the strategy x_{ijk} would have been set to F .
3. In all other cases the Delayer defers to the Prover.

For the analysis we need two real parameters: $q \in (0, 1)$ and $s \in (0, 3 - 3q)$. Since we need $\frac{1}{a} + \frac{1}{b} = 1$ we set $b = \frac{a}{a-1}$.

We now show that this strategy guarantees that the Delayer gets at least Dc points. Since the Delayer will *never* allow a monochromatic rectangle the game ends when there is some i, j such that

$$x_{ij1} = x_{ij2} = \dots = x_{ijc} = F.$$

Who set these variables to F ? Either at least qc were set to F by the Prover or at least $(1 - q)c$ were set to F by the Delayer. This leads to several cases.

1. At least qc were set to F by the Prover. The Delayer gets at least $qc \lg a$ points.
2. At least $(1 - q)c$ were set to F by the Delayer. For every k such that the Delayer set x_{ijk} to F there is an (i', j') (with $i \neq i'$ and $j \neq j'$) such that $x_{i'jk}$, $x_{ij'k}$, and $x_{i'j'k}$ were all set to T (we do not know by who). Consider the variables we know were set to T because Delayer set x_{ijk} to F . These variables all have the last subscript of k . Therefore these sets-of-three variables associated to each x_{ijk} are disjoint. Hence there are at least $3(1 - q)c = (3 - 3q)c$ variables that were set to T . There are two cases. Recall that $s \in (0, 3 - 3q)$.
 - (a) The Prover set at least sc of them to T . Then the Delayer gets at least $sc \lg(b) = sc \lg(a/(a - 1))$ points.
 - (b) The Delayer set at least $(s - (3 - 3q))c = (s + 3q - 3)c$ of them to T . If the Delayer is setting some variable $x_{i'j'k}$ to T it's because the Prover set rc others of the form $x_{i'j'k'}$ to F . These sets-of- rc -variables are all disjoint. Hence the Prover set at least $(s + 3q - 3)rc^2$ variables to F . Therefore the Delayer gets at least $(s + 3q - 3)rc^2 \lg a$ points.

We need to set $a \in (1, \infty)$, $q, r \in (0, 1)$, and $s \in (0, 3 - 3q)$ to maximize the minimum of

1. $qc \lg a$
2. $sc \lg(a/(a - 1))$
3. $(s + 3q - 3)rc^2$

We optimize our choices by setting $qc \lg a = sc \lg(a/(a - 1))$ (approximately) and thinking (correctly) that the c^2 term in $(s + 3q - 3)rc^2$ will force this term to be large when c is large. To achieve this we take

- $q = 0.56415$. Note that $3 - 3q = 1.30755$.
- $s = 1.30754$. Note that $s \in (0, 3 - 3q)$.
- $r = 0.9$. Note that $(s + 3q - 3)r = (0.00001) * 0.9 = 0.00009$. (Any value of $r \in (0, 1)$ would have sufficed.)
- $a = 2.793200$
- $b = a/(a - 1) = 1.557662$ (approximately)

Using these values we get $qc \lg a, sc \lg(a/(a - 1)) \geq 0.836$. We want

$$(0.00009c^2) \geq 0.836c$$

$$(0.00009c) \geq 0.836$$

$$c \geq 9288$$

With this choice of parameters, for $c \geq 9288$, the Delayer gets at least $0.836c$ points. Hence any tree resolution proof of $GRID(n, m, c)$ must have size at least $2^{0.836c}$. ■

6 Lower Bounds on Tree Res for Other Shapes

We did not use any property specific to rectangles in our proof of Theorem 5.5. We can generalize our result to any other shape; however, the constant in $2^{\Omega(c)}$ will change.

First we give a definition of rectangle that will help us to generalize it.

Def 6.1 Let $c, N, M \in \mathbb{N}$. A (full or partial) mapping of $G_{N,M}$ to $\{1, \dots, c\}$ is a c -coloring if there does not exist a set of points $\{(a, b), (a + t, b), (a, b + s), (a + t, b + s)\}$ that are all the same color.

Look at the points

$$\{(a, b), (a + t, b), (a, b + s), (a + t, b + s)\}.$$

We can view them as

$$\{(s \times 0, t \times 0) + (a, b), (s \times 1, t \times 0) + (a, b), (s \times 0, t \times 1) + (a, b), (s \times 1, t \times 1) + (a, b)\}.$$

Informally, the set of rectangles is generated by $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$. Formally we can view the set of rectangles on the lattice points of the plane (upper quadrant) as the intersection of $\mathbb{N} \times \mathbb{N}$ with

$$\bigcup_{s, t, a, b \in \mathbb{Q}} \{ \{ (s \times 0, t \times 0) + (a, b), (s \times 1, t \times 0) + (a, b), (s \times 0, t \times 1) + (a, b), (s \times 1, t \times 1) + (a, b) \} \}$$

Note that the pair of curly braces is not a typo. We are looking at sets of 4-sets of points. We generalize rectangles.

Def 6.2 Let

$$S = \{(x_1, y_1), \dots, (x_L, y_L)\}$$

be a set of lattice points in the plane. Let

$$\text{stretch}(S) = \bigcup_{s, t \in \mathbb{Q}} \{ \{ (sx_1, ty_1), \dots, (sx_L, ty_L) \} \}$$

and

$$\text{translate}(S) = \bigcup_{a, b \in \mathbb{Q}} \{ \{ (x_1 + a, y_1 + b), \dots, (x_L + a, y_L + b) \} \}.$$

These are the sets of points we will be trying to avoid making monochromatic. Hence let

$$\text{avoid}(S) = \text{translate}(\text{stretch}(S)).$$

We can now generalize the rectangle problem.

Def 6.3 Let $N, M \in \mathbb{N}$ and S be a set of lattice points. A (partial or full) mapping χ from $G_{N,M}$ into $[c]$ is a (c, S) -coloring if there are no monochromatic sets in $\text{avoid}(S)$.

Def 6.4 Let $N, M \in \mathbb{N}$ and S be a set of lattice points. Let $GRID(n, m, c; S)$ be the Boolean formula that can be interpreted as saying that $G_{n,m}$ is (c, S) -colorable. We omit details.

The following theorems have proof similar to those in Section 5.

Theorem 6.5 *Let $n, m \in \mathbb{N}$ and S be a set of lattice points. Let n, m, c be such that $G_{n,m}$ is not (c, S) -colorable. Any tree resolution proof of $GRID(n, m, c; S) \notin SAT$ requires size $2^{\Omega(c)}$. The constant in the $\Omega(c)$ depends only on $|S|$ and not the nature of S .*

One could look at other ways to move the points in S around. There is one we find particular interesting. We motivate our definition.

What if we wanted to look at colorings that avoided a monochromatic *square*? The square

$$\{(a, b), (a + s, b), (a, b + s), (a + s, b + s)\}$$

can be viewed as

$$\{(0, 0) + (a, b), (s, 0) + (a, b), (0, s) + (a, b), (s, s) + (a, b)\}.$$

We generalize this.

Def 6.6 Let

$$S = \{(x_1, y_1), \dots, (x_L, y_L)\}$$

be a set of lattice points in the plane. Let

$$\text{halfstretch}(S) = \bigcup_{s \in \mathbb{Q}} \{(sx_1, sy_1), \dots, (sx_L, sy_L)\}$$

These are the sets of points we will be trying to avoid making monochromatic. We would like to call it “avoid” but that name has already been taken; hence we call it avoid_2 .

$$\text{avoid}_2(S) = \text{translate}(\text{halfstretch}(S)).$$

(Note that the 2 has no significance. It is just there to distinguish avoid and avoid_2 .)

Def 6.7 Let $N, M \in \mathbb{N}$ and S be a set of lattice points. A (partial or full) mapping χ from $G_{N,M}$ into $[c]$ is a $(c, S)_2$ -coloring if there are no monochromatic sets in $\text{avoid}_2(S)$. (Note that the 2 has no significance. It is just there to distinguish (c, S) and $(c, S)_2$.)

Def 6.8 Let $N, M \in \mathbb{N}$ and S be a set of lattice points. A (partial or full) mapping χ from $G_{N,M}$ into $[c]$ is a $(c, S)_2$ -coloring if there are no monochromatic sets in $\text{avoid}_2(S)$.

We can now generalize the square problem.

Def 6.9 Let $N, M \in \mathbb{N}$ and S be a set of lattice points. Let $GRID(n, m, c; S_2)$ be the Boolean formula that can be interpreted as saying that $G_{n,m}$ is $(c, S)_2$ -colorable. We omit details.

The following theorems have proof similar to those in Section 5.

Theorem 6.10 Let $n, m \in \mathbb{N}$ and S be a set of lattice points such that $|S| \geq 2$. Let n, m, c be such that $G_{n,m}$ is not $(c, S)_2$ -colorable. Any tree resolution proof of $GRID(n, m, c)_2 \notin SAT$ requires size $2^{\Omega(c)}$.

7 Lower Bound on CP-Tree Res for $GRID(c+1, c\binom{c}{2} + 1, c)$

By Lemma 4.3 the formula $GRID(c+1, c\binom{c}{2} + 1, c) \notin SAT$. Note that it's just barely not satisfiable since $GRID(c+1, c\binom{c}{2}, c) \in SAT$. In this section we show that any Cutting Plane Tree Resolution proof that $GRID(c+1, c\binom{c}{2} + 1, c) \notin SAT$ requires size $2^{\Omega(c^3/\log^2 c)}$.

Notation 7.1 Let A be an integer valued matrix and \vec{b} be an integer valued vector such that there is no 0-1 vector \vec{x} with $A\vec{x} \leq \vec{b}$. We refer to this as $A\vec{x} \leq \vec{b} \notin SAT$.

Any CNF-formula can be phrased in this form with only a linear blowup in size. For every variable x we have variables x and \bar{x} and the inequalities

$$\begin{aligned} x + \bar{x} &\leq 1 \\ -x - \bar{x} &\leq -1 \end{aligned}$$

If C is a clause with literals L_1, \dots, L_k then we have the inequality

$$L_1 + \dots + L_k \geq 1$$

In particular, the formulas $GRID(n, m, c)$ can be put in this form.

7.1 Background on CP-Tree Resolution and Link to Communication Complexity

The definitions of Cutting Plane Proofs and Tree Cutting Plane Proofs are standard. The connection to communication complexity (Lemma 7.5) is from [8] (see also [9] Lemmas 19.7 and 19.11).

Def 7.2 A *Cutting Planes Proof that $A\vec{x} \leq \vec{b} \notin SAT$* (henceforth CP Proof) is a sequence of linear inequalities such that on each line you have either

1. One of the inequalities in $A\vec{x} \leq \vec{b}$ (called an AXIOM).
2. If $\vec{a}_1 \cdot \vec{x} \leq c_1$ and $\vec{a}_2 \cdot \vec{x} \leq c_2$ are on prior lines then $(\vec{a}_1 + \vec{a}_2) \cdot \vec{x} \leq c_1 + c_2$ can be on a line.
3. If $\vec{a} \cdot \vec{x} \leq c$ is on a prior line and $d \in \mathbb{N}$ then $d(\vec{a} \cdot \vec{x}) \leq dc$ can be on a line. (Also if $d \in \mathbb{Z} - \mathbb{N}$ then reverse the inequality.)
4. If $c(\vec{a} \cdot \vec{x}) \leq d$ is on a prior line then $\vec{a} \cdot \vec{x} \leq \lfloor \frac{d}{c} \rfloor$ can be on a line.
5. The last line is an arithmetically false statement (e.g., $1 \leq 0$).

It is easy to see that if there is a cutting planes proof that $A\vec{x} \leq \vec{b} \notin SAT$ then indeed $A\vec{x} \leq \vec{b} \notin SAT$. The converse is also true though slightly harder to prove.

Def 7.3 A *Tree-like CP proof* is one whose underlying graph is a tree.

Def 7.4 Let A be an integer valued matrix and \vec{b} be an integer valued vector such that $A\vec{x} \leq \vec{b} \notin SAT$. Let P_1, P_2 be a partition of the variables in \vec{x} . The Communication Complexity problem $FI(A, \vec{b}, P_1, P_2)$ is as follows.

1. For every variable in P_1 Alice is given a value (0 or 1).
2. For every variable in P_2 Bob is given a value (0 or 1).
3. These assignments constitute an assignment to all of the variables which we denote \vec{x} .
4. Alice and Bob need to determine an inequality in $A\vec{x} \leq \vec{b}$ that is not true.

Lemma 7.5 Let A be an integer valued matrix and \vec{b} be an integer valued vector such that $A\vec{x} \leq \vec{b} \notin SAT$. Let n be the number of variables in \vec{x} . If there is a partition P_1, P_2 of the variables such that, for all ϵ , $R_\epsilon(FI(A, \vec{b}, P_1, P_2)) = \Omega(t)$ then any tree-like CP proof of $A\vec{x} \leq \vec{b}$ requires size $2^{\Omega(t/\log^2 n)}$.

7.2 Lemmas on Communication Complexity

Def 7.6

1. The *Hamming weight* of a binary string x , denoted $w(x)$, is the number of 1's in x .
2. The *Hamming distance* between two, equal-length, binary strings x and y , denoted $d(x, y)$, is the number of positions in which they differ.
3. For a communication problem \mathcal{P} , $D(\mathcal{P})$ denotes the deterministic communication complexity of \mathcal{P} and $R_\epsilon(\mathcal{P})$ denotes the randomized public coin communication complexity of \mathcal{P} with error $\leq \epsilon$.

Def 7.7 We define several communication complexity problems.

1. PHPstr_n : Alice gets a string $x \in \Sigma^n$, and Bob gets a string $y \in \Sigma^n$ with $|\Sigma| = 2n - 1$. They are promised that for all $i \neq j$, the letters x_i and x_j (resp. y_i and y_j) are distinct. By the PHP, there must exist at least one $(i, j) \in [n] \times [n]$ such that $x_i = y_j$. They are further promised that (i, j) is *unique*. *The goal is to find (i, j)* . (Alice learns i , and Bob learns j .)
2. PHPset_n : Alice gets a set $x \in \binom{\Sigma}{n}$, and Bob gets a set $y \in \binom{\Sigma}{n}$ with $|\Sigma| = 2n - 1$. By the PHP, there must exist at least one $\sigma \in \Sigma$ such that $\sigma \in x \cap y$. They are further promised that σ is *unique*. *The goal is to find σ* . (Both learn σ .)
3. PrMeet_n : Alice gets a string $x \in \{0, 1\}^n$, and Bob gets a string $y \in \{0, 1\}^n$ with $n = 2m - 1$. They are promised that (1) $w(x) = w(y) = m$, that (2) there is a *unique* $i \in [n]$ such that $x_i = y_i = 1$, and that (3) for all $j \neq i$, $(x_j, y_j) \in \{(0, 1), (1, 0)\}$. *The goal is to find i* . (Both learn i .)
4. UM_n : (called the *universal monotone relation*) Alice is given $x \in \{0, 1\}^n$, and Bob is given $y \in \{0, 1\}^n$. They are promised that there exists i such that $x_i = 1$ and $y_i = 0$. *The goal is to find some such i* . (Both learn i .)
5. PrUM_n : This is a restriction of UM_n . They are additionally promised (1) $n = 2m - 1$ is odd, (2) $w(x) = m$, (3) $w(y) = m - 1$, and (4) $d(x, y) = 1$. Hence (a) there is a *unique* index $i \in [n]$ such that $x_i = 1$ and $y_i = 0$, (b) for all $j \neq i$, $(x_j, y_j) \in \{(0, 0), (1, 1)\}$, and moreover (c) these $(0, 0)$'s and $(1, 1)$'s occur in an equal number. *The goal is to find i* . (Both learn i .)
6. DISJ_n : Alice gets a string $x \in \{0, 1\}^n$, and Bob gets a string $y \in \{0, 1\}^n$. They need to *decide* if x and y intersect ($\exists i$ where $x_i = y_i$).
7. PrDISJ_n : $n = 2m + 1$ is odd. Alice gets a string $x \in \{0, 1\}^n$, and Bob gets a string $y \in \{0, 1\}^n$. They are promised that x and y have exactly $m + 1$ 1's and m 0's and intersect at most once. They need to *decide* if x and y intersect ($\exists i$ where $x_i = y_i$).

We will need the following notion of reduction.

Def 7.8 Let f, g be a communication problem. It can be a decision, a function, and/or a promise problem.

1. $f \leq_{cc} g$ if there exists a protocol for f that has the following properties.
 - (a) The protocol may invoke a protocol for g once on an input of length $O(n)$.
 - (b) Before and after the invocation, the players may communicate polylog bits.

The following lemma is obvious.

Lemma 7.9 *If $f \leq_{cc} g$ and $(\forall \epsilon)[R_\epsilon(f) = \Omega(n)]$ then $(\forall \epsilon)[R_\epsilon(g) = \Omega(n)]$.*

Lemma 7.10 *For all ϵ $R_\epsilon(\text{PrUM}_n) = \Omega(n)$.*

Proof: In [11] it was shown that $\text{DISJ}_n \leq_{cc} \text{UM}_n$. A closer examination of the proof shows that it also shows $\text{PrDISJ}_n \leq_{cc} \text{PrUM}_n$.

Kalyanasundaram and Schnitger [10] showed that, for all ϵ , $R_\epsilon(\text{DISJ}_n) = \Omega(n)$. Razborov [14] has a simpler proof where he only looks at inputs that satisfy the promise of PrDISJ_n . Hence he showed $R_\epsilon(\text{PrDISJ}_n) = \Omega(n)$. From $\text{PrDISJ}_n \leq_{cc} \text{PrUM}_n$, $R_\epsilon(\text{PrDISJ}_n) = \Omega(n)$, and Lemma 7.9 the result follows. ■

Lemma 7.11

1.

$$\text{PrUM}_n \leq_{cc} \text{PrMeet}_n \leq_{cc} \text{PHPset}_{(n+1)/2}.$$

(The last reduction only holds when n is odd.)

2.

$$\text{PHPset}_n \leq_{cc} \text{PHPstr}_n.$$

3. For all ϵ $R_\epsilon(\text{PHPstr}_n) = \Omega(n)$. (This follows from parts 1,2 and Lemmas 7.9, 7.10.)

Proof:

$\text{PrUM}_n \leq_{cc} \text{PrMeet}_n$: Alice gets $x \in \{0, 1\}^n$, Bob gets $y \in \{0, 1\}^n$ so that (x, y) satisfies the promise of PrUM_n . Let $n = 2m - 1$. We show that (x, \bar{y}) satisfies the promise of PrMeet_n and that $\text{PrUM}_n(x, y) = \text{PrMeet}_n(x, \bar{y})$.

Since $w(y) = m - 1$, $w(\bar{y}) = n - (m - 1) = m$. We still have $w(x) = m$ so $w(x) = w(y) = m$. Since there is a unique $i \in [n]$ such that $x_i = 1$ and $y_i = 0$, then must be a unique $i \in [n]$ (the same one) such that $x_i = \bar{y}_i = 1$. (This establishes $\text{PrUM}_n(x, y) = \text{PrMeet}_n(x, \bar{y})$.) Since for all all $j \neq i$, $(x_j, y_j) \in \{(0, 0), (1, 1)\}$, for all $j \neq i$, $(x_j, \bar{y}_j) \in \{(0, 1), (1, 0)\}$.

$\text{PrMeet}_n \leq_{\text{cc}} \text{PHPset}_{(n+1)/2}$: Alice gets $x \in \{0,1\}^n$, Bob gets $y \in \{0,1\}^n$ so that (x,y) satisfies the promise of PrMeet_n . Let $m = (n+1)/2$. Note that $w(x) = w(y) = m$. Let Σ be an alphabet of size n . Both Alice and Bob agree on an ordering of Σ ahead of time.

Alice views her n -bit string x (resp. Bob views his string y) as the bit vector of an m -element subset of Σ . We denote this subset a (and for Bob b). Clearly (a,b) satisfies the promise of $\text{PHPset}_{(n+1)/2}$ and $\text{PrMeet}_n(x,y) = \text{PHPset}_{(n+1)/2}(a,b)$.

$\text{PHPset}_n \leq_{\text{cc}} \text{PHPstr}_n$: Σ is an alphabet of size $2n-1$. Alice and Bob agree on an ordering of Σ ahead of time. Alice gets $x \in \binom{\Sigma}{n}$, Bob gets $y \in \binom{\Sigma}{n}$. The sets x,y satisfy the promise of PHPset_n .

Alice (Bob) forms the string $x' \in \Sigma^n$ ($y' \in \Sigma^n$) which is the elements of x (y) written in order. Clearly x', y' satisfy the promise of PHPstr_n . Alice and Bob run the protocol for PHPstr_n on (x', y') . Alice obtains i , Bob obtains j . The i th element of x' is the same as the j th element of y' . This element is σ which is promised in $\text{PHPset}_n(x,y)$. ■

7.3 Lower Bound on CP-Tree Resolution for $\text{GRID}(c+1, c\binom{c}{2}+1, c)$

Theorem 7.12 *Let $A\vec{x} \leq \vec{b}$ be the translation of $\text{GRID}(c+1, c\binom{c}{2}+1, c)$ into an integer program. Any Tree-CP proof that $A\vec{x} \leq \vec{b} \notin \text{SAT}$ requires $2^{\Omega(c^3/\log^2 c)}$ size.*

Proof: We do the case where $c\binom{c}{2}+1$ is even (so $c \equiv 3 \pmod{4}$). The other cases are similar but require slight variants of Lemma 7.11.

Split the $(c+1) \times c\binom{c}{2}+1$ evenly into two halves, both of size $((c+1) \times c\binom{c}{2}+1)/2$. Let P_1, P_2 be the partition of the variables so that Alice gets all of the variables involved in coloring the left half, and Bob gets all of the variables involved in coloring the right half. We show that $D(CC(A, \vec{b}, P_1, P_2)) = \Omega(c^3)$. Note that the number of variables is $\Theta(c^4)$. Hence, by Lemma 7.5 we obtain that the size of any Tree-CP proof of $A\vec{x} \leq \vec{b} \notin \text{SAT}$ requires size $2^{\Omega(c^3/\log^2 c)}$.

We restrict the problem to the case where every column has $c-1$ colors occurring once and the remaining color occurring twice. Hence one can view a coloring as string of length $2m = c\binom{c}{2}+1$ over an alphabet of size $n = c\binom{c}{2}$. Note that Alice and Bob each get a string of length m over an alphabet of size $n = 2m-1$.

In order to find which inequality is violated Alice and Bob need to find which columnn they agree on (e.g., Alice's column i is the same as Bob's column j). This is precisely the problem PHPstr_n . Hence, by Lemma 7.11 this problem has communication complexity $\Omega(n) = \Omega(c^3)$. Therefore, by Lemma 7.5, any Tree-CP proof of $A\vec{x} \leq \vec{b}$ requires $2^{c^3/\log^2 c}$. ■

Lower bounds on Tree-CP proofs yield lower bounds on Tree-Resolution (with a constant factor loss) (see Prop 19.4 of [9]). Hence we have the following.

Corollary 7.13 *Any Tree-resolution proof of $\text{GRID}(c+1, c\binom{c}{2}+1, c) \notin \text{SAT}$ requires $2^{\Omega(c^3/\log^2 c)}$ size.*

8 Open Problems

8.1 Open Problems Related to NP-Completeness

Open Problem 1: For which sets of lattice points S is the following problem NP-complete?

$$\{(N, M, c, \chi) \mid \chi \text{ can be extended to a } (c, S)\text{-coloring of } G_{N,M}\}$$

Open Problem 2: For which sets of lattice points S is the following problem NP-complete?

$$\{(N, M, c, \chi) \mid \chi \text{ can be extended to a } (c, S)_2\text{-coloring of } G_{N,M}\}$$

Open Problem 3: Improve our FPT algorithm. Develop an FPT algorithm for the variants we have discussed.

Open Problem 4: Prove that grid coloring problems starting with the empty grid are hard. This may need a new formalism.

8.2 Open Problems Related to Lower Bounds on Tree Resolution

If ϕ is a Boolean formula on v variables then it has a Tree Resolution proof of size $2^{O(v)}$. Hence there is a tree resolution proof of $GRID(n, m, c)$ of size $2^{O(nmc)}$. For particular values of m, n (functions of c) can we do better? We have already obtained this kind of result for $GRID(c + 1, c \binom{c}{2} + 1, c)$ (see Corollary 7.13).

Open Problem 1: For various n and m that are functions of c such that $G_{n,m}$ is not c -colorable, obtain a better lower bound on Tree Resolution than $2^{\Omega(c)}$.

There are unsatisfiable Boolean formulas for which Tree Resolution requires exponential size, but there are polynomial size resolution proofs.

Open Problem 2: Determine upper and lower bounds for the size of Resolution proofs of $GRID(n, m, c)$.

8.3 Open Problems Related to Lower Bounds on Tree-CP Refutations

We showed that Tree-CP for $GRID(c + 1, c \binom{c}{2} + 1, c) \notin SAT$ require exponential size. For other families of non- c -colorable grids either show that tree CP proof requires exponential size or show that there are short tree CP proofs. For other families of non- c -colorable grids either show that (non-tree) CP proofs requires exponential size or show that CP proofs are short.

9 Acknowledgments

We would like to thank Clyde Kruskal for proofreading and discussion. We would like to thank Stasys Jukna whose marvelous exposition of the Prover-Delayer games and the tree-CP proofs inspired the second and third parts of this paper, and for some technical advice (translation: he found a bug and helped us fix it). We would like to thank Daniel Marx for pointing out an improvement in the fixed parameter algorithm which we subsequently used. We would like to thank Wing-ning Li for pointing out that the case of n, m binary, while it seems to not be in NP, is actually unknown. We would also like to thank Tucker Bane, Richard Chang, Peter Fontana, David Harris, Jared Marx-Kuo, Jessica Shi, and Marius Zimand, for listening to Bill present these results and hence clarifying them.

References

- [1] O. Beyersdorr, N. Galesi, and M. Lauria. A lower bound for the pigeon-hole principle in the tree-like resolution asymmetric prover-delayer games. *Information Processing Letters*, 110, 2010. The paper and a talk on it are here: <http://www.cs.umd.edu/~gasarch/resolution.html>.
- [2] S. Fenner, W. Gasarch, C. Glover, and S. Purewal. Rectangle free colorings of grids, 2012. <http://arxiv.org/abs/1005.3750>.
- [3] W. Gasarch. The 17×17 challenge. Worth \$289.00. This is not a joke, 2009. <http://blog.computationalcomplexity.org/2009/11/17x17-challenge-worth-28900-this-is->
- [4] W. Gasarch. A possible NP-intermediary problem. <http://blog.computationalcomplexity.org/2010/04/possible-np-intermediary-problem.htm> 2010.
- [5] W. Gasarch. The 17×17 SOLVED! (also 18×18) <http://blog.computationalcomplexity.org/2012/02/17x17-problem-solved-also-18x18.htm> 2012.
- [6] B. Hayes. The 17×17 challenge, 2009. <http://bit-player.org/2009/the-17x17-challenge>.
- [7] S. Homer and L. Longpre. On reductions of NP sets to sparse sets. *Journal of Computer and System Sciences*, 48, 1994. Prior version in *STRUCTURES* 1991.
- [8] R. Impagliazzo and T. P. and. Upper and lower bounds for tree-like cutting planes proofs. In *Proceedings of the Ninth Annual IEEE Symposium on Logic in Computer Science*, Paris, France, 1994. <http://www.cs.toronto.edu/~toni>.
- [9] S. Jukna. *Boolean function complexity: advances and frontiers*. Algorithms and Combinatorics Vol 27. Springer, 2012.

- [10] B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5:545–557, 1992. Prior version in *Conf. on Structure in Complexity Theory*, 1987 (STRUCTURES).
- [11] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [12] S. Mahaney. Sparse complete sets for NP: Solution to a conjecture of Berman and Hartmanis. *Journal of Computer and System Sciences*, 25:130–143, 1982.
- [13] P. Pudlka and R. Impagliazzo. A lower bound for DLL algorithms for SAT. In *Eleventh Symposium on Discrete Algorithms: Proceedings of SODA '00*, 2000.
- [14] A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106:385–390, 1992. Prior version in ICALP 1990. Available online at <http://people.cs.chicago.edu/~razborov/research>.
- [15] B. Steinbach and C. Posthoff. Extremely complex 4-colored rectangle-free grids: Solution of an open multiple-valued problem. In *Proceedings of the Forty-Second IEEE International Symposia on Multiple-Valued Logic*, 2012. <http://www.cs.umd.edu/~gasarch/PAPERSR/17solved.pdf>.